

Transport Area Working Group
(tsvwg)
Internet-Draft
Updates: [792](#), [1122](#), [1812](#)
(if approved)
Intended status: Standards Track
Expires: June 9, 2011

F. Gont
UTN/FRH
December 6, 2010

Deprecation of ICMP Source Quench messages
draft-gont-tsvwg-source-quench-01.txt

Abstract

This document formally deprecates the use of ICMP Source Quench messages by transport protocols.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on June 9, 2011.

Copyright Notice

Copyright (c) 2010 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	3
2.	ICMP Source Quench messages	3
3.	Updating RFC 1122	4
4.	Updating RFC 1812	4
5.	General Advice to Transport Protocols	4
6.	Changing the status of RFC 1016 Historic	4
7.	Security Considerations	4
8.	IANA Considerations	5
9.	Acknowledgements	5
10.	References	5
10.1.	Normative References	5
10.2.	Informative References	6
Appendix A.	Survey of support of ICMP Source Quench in some popular TCP/IP implementations	6
Appendix B.	Changes from previous versions of the draft (to be removed by the RFC Editor before publishing this document as an RFC)	6
B.1.	Changes from draft-gont-tsvwg-source-quench-00	6
Author's Address	7

1. Introduction

The ICMP specification [[RFC0792](#)] defines the ICMPv4 Source Quench message (type 4, code 0), which is meant as a mechanism for congestion control. ICMP Source Quench is known to be an ineffective (and unfair) antidote for congestion, and generation of ICMP Source Quench messages by routers has been deprecated by [[RFC1812](#)] for a long time. However, reaction to ICMP Source Quench messages in transport protocols has never been formally deprecated.

This document formally deprecates reaction to ICMP Source Quench messages by transport protocols such as TCP.

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC 2119](#) [[RFC2119](#)].

2. ICMP Source Quench messages

The ICMP specification [[RFC0792](#)] defines the ICMP Source Quench message (type 4, code 0), which is meant to provide a mechanism for congestion control. The Host Requirements RFC [[RFC1122](#)] states in [Section 4.2.3.9](#) that hosts MUST react to ICMP Source Quench messages by slowing transmission on the connection, and further adds that the RECOMMENDED procedure is to put the corresponding connection in the slow-start phase of TCP's congestion control algorithm [[RFC5681](#)].

[[RFC1812](#)] notes that research suggests that ICMP Source Quench is an ineffective (and unfair) antidote for congestion, and formally deprecates the generation of ICMP Source Quench messages by routers, stating that routers SHOULD NOT send ICMP Source Quench messages in response to congestion.

[[RFC5927](#)] discusses the use of ICMP Source Quench messages for performing "blind throughput-reduction" attacks, and notes that most TCP implementations silently ignore ICMP Source Quench messages.

We note that TCP implements its own congestion control mechanisms [[RFC5681](#)] [[RFC3168](#)], that do not depend on ICMP Source Quench messages.

It is interesting to note that ICMPv6 [[RFC4443](#)] does not specify a "Source Quench" message.

Gont

Expires June 9, 2011

[Page 3]

3. Updating [RFC 1122](#)

We hereby update [Section 3.2.2.3 of \[RFC1122\]](#) as follows:

A host SHOULD NOT send ICMP Source Quench messages.

If Source Quench message is received, the IP layer MAY silently discard it.

[Section 4.2.3.9 of \[RFC1122\]](#) is updated as follows:

TCP SHOULD silently discard any received ICMP Source Quench messages.

4. Updating [RFC 1812](#)

We hereby update [Section 4.3.3.3 OF \[RFC1812\]](#) as follows:

A router SHOULD ignore any ICMP Source Quench messages it receives.

5. General Advice to Transport Protocols

If an ICMP Source Quench message is received by a transport-protocol instance (e.g., a TCP connection), it SHOULD be silently ignored.

6. Changing the status of [RFC 1016](#) Historic

This document requests the RFC Editor to change the status of [\[RFC1016\]](#) to "Historic".

7. Security Considerations

ICMP Source Quench messages could be leveraged for performing blind throughput-reduction attacks against TCP and similar protocols. This attack vector, along with possible countermeasures, have been discussed in great detail in [\[RFC5927\]](#) and [\[CPNI-TCP\]](#). However, as noted in [\[RFC5927\]](#) and [\[CPNI-TCP\]](#), virtually all current versions of popular TCP implementations already silently ignore ICMP Source Quench messages.

Silently ignoring ICMP Source Quench messages, as specified in this document, eliminates the aforementioned attack vector.

If deemed necessary, ICMP Source Quench messages could be filtered at firewalls.

8. IANA Considerations

This document has no actions for IANA. The RFC-Editor can remove this section before publication of this document as an RFC.

9. Acknowledgements

The author of this document would like to thank (in alphabetical order) Fred Baker, Scott Bradner, Antonio De Simone, Gorrry Fairhurst, Alfred Hoenes, Mahesh Jethanandani, Dan Wing, and Andrew Yourtchenko, for providing valuable feedback on earlier versions of this document.

This document has benefited from discussions within the TCCM Working Group while working on [[RFC5927](#)].

10. References

10.1. Normative References

- [RFC0792] Postel, J., "Internet Control Message Protocol", STD 5, [RFC 792](#), September 1981.
- [RFC0793] Postel, J., "Transmission Control Protocol", STD 7, [RFC 793](#), September 1981.
- [RFC1016] Prue, W. and J. Postel, "Something a host could do with source quench: The Source Quench Introduced Delay (SQuID)", [RFC 1016](#), July 1987.
- [RFC1122] Braden, R., "Requirements for Internet Hosts - Communication Layers", STD 3, [RFC 1122](#), October 1989.
- [RFC1812] Baker, F., "Requirements for IP Version 4 Routers", [RFC 1812](#), June 1995.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC3168] Ramakrishnan, K., Floyd, S., and D. Black, "The Addition of Explicit Congestion Notification (ECN) to IP", [RFC 3168](#), September 2001.

- [RFC4443] Conta, A., Deering, S., and M. Gupta, "Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification", [RFC 4443](#), March 2006.
- [RFC5681] Allman, M., Paxson, V., and E. Blanton, "TCP Congestion Control", [RFC 5681](#), September 2009.

10.2. Informative References

- [CPNI-TCP] CPNI, "Security Assessment of the Transmission Control Protocol (TCP)", <http://www.cpni.gov.uk/Docs/tn-03-09-security-assessment-TCP.pdf>, 2009.
- [FreeBSD] The FreeBSD Project, "<http://www.freebsd.org>".
- [Linux] The Linux Project, "<http://www.kernel.org>".
- [NetBSD] The NetBSD Project, "<http://www.netbsd.org>".
- [OpenBSD] The OpenBSD Project, "<http://www.openbsd.org>".
- [RFC5927] Gont, F., "ICMP Attacks against TCP", [RFC 5927](#), July 2010.

Appendix A. Survey of support of ICMP Source Quench in some popular TCP/IP implementations

A large number of implementations completely ignore ICMP Source Quench messages meant for TCP connections. This behavior has been implemented in, at least, Linux [[Linux](#)] since 2004, and in FreeBSD [[FreeBSD](#)], NetBSD [[NetBSD](#)], and OpenBSD [[OpenBSD](#)] since 2005.

Appendix B. Changes from previous versions of the draft (to be removed by the RFC Editor before publishing this document as an RFC)

B.1. Changes from [draft-gont-tsvwg-source-quench-00](#)

- o This revision reflects the recent discussion about ICMP Source Quench messages on the tsvwg mailing-list. A detailed list of the changes is available at:
<http://www.ietf.org/mail-archive/web/tsvwg/current/msg10407.html>

Author's Address

Fernando Gont
Universidad Tecnologica Nacional / Facultad Regional Haedo
Evaristo Carriego 2644
Haedo, Provincia de Buenos Aires 1706
Argentina

Phone: +54 11 4650 8472
Email: fernando@gont.com.ar
URI: <http://www.gont.com.ar>